AIRCURRENTS

PORTFOLIO OPTIMIZATION FOR INSURANCE COMPANIES

01.2011

Editor's note: AIR recently launched a decision analytics division within its consulting and client services group. Its offerings include novel solutions for portfolio optimization that are based on evolutionary search algorithms. In this article, Dr. Guillermo Franco, Manager and Principal Engineer, discusses methods applicable to insurers.

By Dr. Guillermo Franco

INTRODUCTION

Portfolio optimization is a familiar concept in the business of insurance, although the term is often used to describe several slightly different processes. For the purposes of this article, we will define it as the selection of a set of policies that maximize certain desirable metrics of performance of a primary insurance book of business, while constraining other undesirable metrics. An insurance company, for example, may want to acquire policies from a state wind-pool or Fair Access to Insurance Requirements (FAIR) plan that satisfy its premium requirements while simultaneously producing the minimum expected loss. During renewal season, another company may want to examine whether dropping some policies could keep premiums at a satisfactory level but decrease their reinsurance costs significantly. In another example, a company may try to minimize expected losses in order to comply with regulatory demands while keeping the risk-based return of the portfolio as high as possible. In a more complex scenario, a company may want to impose certain constraints on the construction types and geographical aggregates included in the portfolio while keeping their financial yields constant.

Using AIR catastrophe models, it is possible to compute complex loss metrics of a portfolio as a function of the location of the exposure or of its physical features. It is therefore mathematically possible to tailor a portfolio so that its numerical performance meets certain desired objectives within reason. From such an analysis, underwriting patterns can be derived in order to guide the company's operations towards achieving its risk management and financial goals.

In the practical realm, solutions readily available in the market to carry out this type of exercise are generally limited to ranking policies using a simple metric (like risk-based return or average annual loss) or, in more sophisticated attempts, incorporating mathematical approaches like the steepest ascent method for finding the maximum value of a function. These methods are typically driven by only one metric and are not particularly well-suited to handle multiple simultaneous constraints. More often than not, these exercises are carried out manually using a type of scenario testing called marginal impact analysis, which involves examining the change to some given metrics upon discarding or accepting a certain portion of the portfolio.

While they may be referred to as such, these types of analyses are not true portfolio optimization exercises, which are rarely conducted for two primary reasons. First, the computation of many risk metrics is nonlinear, meaning that these metrics need to be recomputed from scratch for each alternative policy combination. Second, the process of solving for the optimal portfolio quickly becomes unfeasible from a technological standpoint because of the staggering number of potential solutions, even with relatively few policies.

Evolutionary search techniques, which are optimization methods that partly rely on the semi-random exploration of the solution space, have been suggested in the academic literature as plausible candidates to solve the problem of complex searches of potential policy combinations in an automated fashion. AIR has devoted several years to the development and application of these optimization techniques for addressing sampling variability constraints in our stochastic catalogs and in producing finely calibrated parametric catastrophe bonds. This article will explain how these techniques are now helping us approach the problem of portfolio optimization from a fresh and promising angle.

A SIMPLE EXAMPLE

Suppose that a portfolio consists of two policies exposed to U.S. hurricane risk. Using the AIR hurricane model, we can calculate losses to each policy for every year of AIR's 10,000-year stochastic catalog. The resulting information can be arranged in a 2 × 10,000 table, with each of the two rows corresponding to a policy, and each column corresponding to one year of simulated hurricane activity.

To obtain the loss exceedance probability (EP) curve for each policy, the losses in each row are sorted from largest to smallest. The largest loss can be assigned an exceedance probability of 0.01% (1/10,000), meaning that it is equaled or surpassed only once in the 10,000 simulation years. The second largest loss is equaled or surpassed twice, corresponding to an exceedance probability of 2 in 10,000 years, or 0.02%. Likewise, the twentieth, fortieth, and hundredth largest losses have exceedance probabilities of 0.2%, 0.4%, and 1%, corresponding to the 500-, 250-, and 100-year return periods, respectively. By assigning a probability of exceedance to each loss, a full EP curve can be derived for each policy.

What about the EP curve for the entire portfolio? Keeping in mind that the loss to the portfolio for a given year is equal to the loss to Policy 1 for that year plus the loss to Policy 2, we must first go back to sum the losses on a year-by-year basis and then re-rank the losses to the portfolio in order to compute a new EP curve. It is not possible to simply add the EP curves of the two policies because at each exceedance probability, the corresponding year of the catalog for Policy 1 may not match that of Policy 2. This non-additive property of the EP curve calculation is a characteristic nonlinearity of risk metrics, and here we encounter the first mathematical complexity in the analysis of portfolio losses.

To determine the impact of the addition or removal of a particular policy on the expected losses of a portfolio, it is necessary to first calculate the total portfolio losses with and without that policy, and then re-compute the EP curve to examine the changes in loss levels and return periods. This type of analysis—known as marginal impact analysis—is routinely conducted using AIR's CLASIC/2™ software.

For example, suppose it is renewal season and we want to determine which policies should be renewed in order to achieve certain business objectives. For our two-policy portfolio, the marginal impact analysis will be quite simple, as there are only four possibilities for the construction of the portfolio: renew neither policy, discard Policy 1 and renew Policy 2, renew Policy 1 and discard Policy 2, or renew both policies. Let's assume that we want to evaluate the resulting portfolio in terms of two simple metrics: the total premium and the total expected loss at 1% probability (100-year return period). The four choices can be labeled as 00, 01, 10, and 11, with the first digit representing the first policy and the second digit representing the second policy. The symbol "0" represents non-renewal and the symbol "1" represents renewal of the respective policy.

After calculating the EP curves of the four choices, the premium and loss metrics can be represented in a graph, as shown in Figure 1. Note that solution 00 yields no premium and no loss since the portfolio is empty. Increasing the number of policies increases the premium, but also the loss. Note that in this example, renewing both policies results in a more efficient growth of premium (as indicated by the shallower slope from 00 to 11) than renewing just one policy. This is not always the case, and it is not possible to foresee whether this will be true for a given pair of policies before carrying out the actual computation.

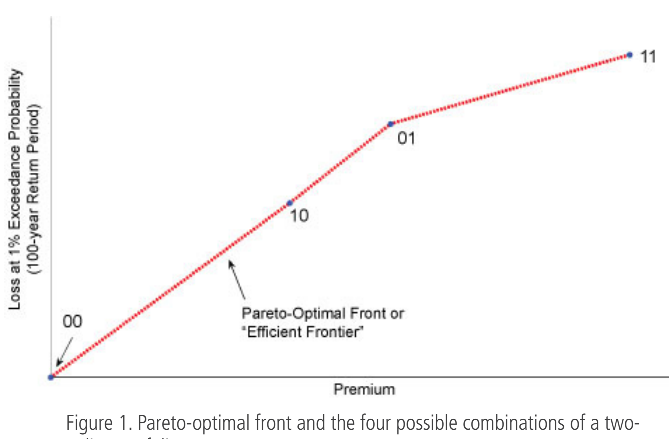

Figure 1. Pareto-optimal front and the four possible combinations of a two-policy portfolio.

The red line that links the solutions is often called the Pareto-optimal front or efficient frontier, and it represents the set of solutions that cannot be judged, *a priori*, explicitly better than one another. That is, in Figure 1, solutions with larger losses also have larger premiums; therefore it is impossible to decide from a mathematical standpoint whether one is superior to another without imposing other criteria. For example, if we wanted to achieve a certain minimum premium or to restrict expected losses below a certain threshold, some solutions may then be revealed to be preferable to others.

The fact that all solutions in this example lie on the Pareto-optimal front is a result of this specific set of exposures and modeling results. It is not always the case, as shown in the next section.

MORE COMPLEX SCENARIOS

To make the problem slightly more complex, assume that the portfolio now contains five policies. As before, we can calculate the expected losses for each year of simulated activity and for each policy, and then compute the EP curve for any combination of policies. With five policies, however, the number of possible combinations grows quite a bit—to 2^5, or 32. Also as before, the different solutions can be labeled with a binary number and a premium-loss graph can be produced, as shown in Figure 2.

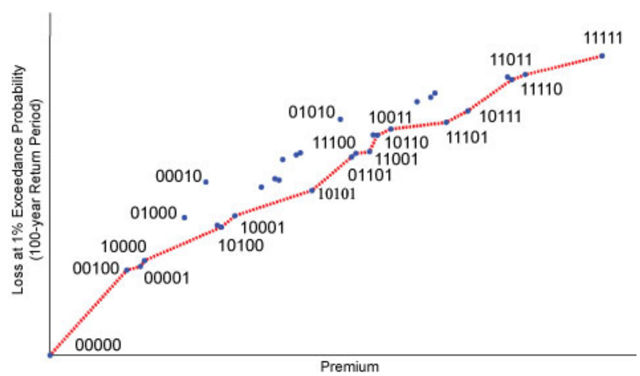

Figure 2. Pareto-optimal front and the 32 possible combinations of a five-policy portfolio.

In this example, many solutions fall outside the Pareto-optimal front.These are mathematically inferior because there are alternative solutions with both lower expected loss and higher premium. There is no logical reason a solution outside the Pareto-optimal front would be selected, based only on the premium and loss criteria, because better performance can be achieved in both metrics. The main problem, however, is to identify those solutions that build the Pareto-optimal front. Which ones are they?

Marginal impact analysis usually leverages some additional information or hypotheses to avoid computing all the possible combinations of policies. For instance, if a company is not able to drop a large portion of the portfolio, many solutions can be eliminated and the number of actual candidates drops significantly.

Sometimes, there is little a priori knowledge to simplify the problem, and although it would be preferable to explore all possible solutions, marginal impact analysis is used, despite its known limitations, because computing all the potential combinations becomes intractable. Consider a portfolio with 100 policies. At first glance, this may not seem like such a challenge, but the number of possible combinations within such a portfolio rises to 2^{100}, or 1.26×10^{30}.

Such a number may be hard to grasp conceptually. You might recall the legend of the invention of chess, for example, in which the king asked the inventor what he wanted as a reward for creating such an ingenious game. The inventor replied: "one grain of wheat for the first square of the chessboard, double as much for the second, double as much for the next, and so on." The king thought that such a trivial reward offended his generosity, but agreed to pay. Certain versions of the legend say that after discovering that the number of wheat grains for the last square totaled 2^{63} (which would weigh over 450 billion tons, some 700 times greater than the present-day annual global wheat production), the king opted instead to chop off the inventor's head.

Returning back to our portfolio optimization problem, a loss table—with the dimensions 100 × 10,000—can be computed as before. To calculate one EP curve for any given portfolio combination using an ordinary desktop computer only takes fractions of a second (about 0.00279 seconds on my computer at AIR). Multiplying this by the number of possible solutions yields the total time needed to exhaustively analyze all the possible choices, which is $2^{100} \times 0.00279$ seconds, approximately 3.54×10^{27} seconds or 1.12×10^{20} years. As this is about 25 billion times the age of our planet (estimated at 4.5 billion years), it becomes apparent that the computational demands in performing portfolio optimization grow enormously for only a modest number of policies. Since most portfolios often contain hundreds or thousands of policies, the problem of exhaustively searching all combinations can be considered mathematically intractable.

GENETIC ALGORITHMS

In the 1970s, mathematician and computer scientist John Holland pioneered the use of *genetic algorithms* to solve search and optimizations problems. Borrowing concepts from Darwin's theory of evolution and Mendel's work on heredity, genetic algorithms are search techniques that rely on randomly driven explorations of the optimization space and are much like the mechanisms behind natural selection—whereby individuals of a given species that possess traits that enhance their chances of survival are more likely to pass down their genetic material to the next generation.

In developing genetic algorithms, Holland applied the fundamental steps that occur during evolution (selection, crossover, and mutation) to a computerized process.[1] First, from a set of randomly generated candidate solutions, a well-performing subset of solutions is selected. During crossover, a new generation of solutions is created by combining the "genetic code" (represented using a binary string, as in previous examples) of different solutions. Often, this is done by mixing a fraction of one solution with a fraction of another to try to preserve the well-performing traits of the parent generation. In addition, some random mutations are introduced into the genetic code by switching several 0s into 1s or vice versa. Mutation rates are kept very low to prevent the destruction of well-performing groups of "genes" (policy selections) in the potential solutions. This process, shown schematically in Figure 3, is reiterated several times, typically with each successive generation of solutions performing better than the previous.

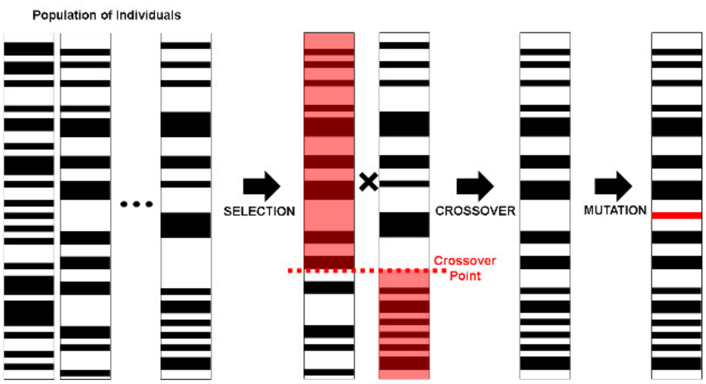

Figure 3. Fundamental operations in genetic algorithms: selection, crossover, and mutation.

Applying genetic algorithms to the portfolio optimization problem, assume that in a portfolio of 10 policies, 20 potential solutions were created at random. Out of these, two of the best performing ones—say 0001001010 and 0110001110—are combined at a random point in the binary string—say after the fourth policy. The crossover operation takes the first four digits from the first solution and combines it with the remaining six digits from the second, resulting in 0001001110. The mutation operation then switches a few genes at random—say in this case, the tenth policy is switched from 0 (reject) to 1 (accept). The final solution after the genetic construction process is 0001001111, which represents a portfolio that contains Policies 4, 7, 8, 9 and 10. This process, shown graphically in Figure 4, is done for other combinations of well-performing portfolios to create a new generation of solutions. This new generation of solutions is then tested and the selection, crossover, and mutation process is reiterated until the best solution is found or until some performance criteria are met.

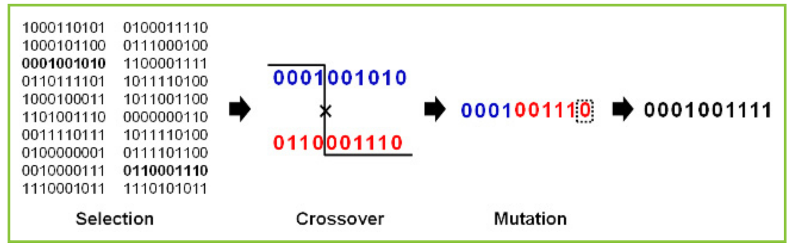

Figure 4. Production of a new solution through the simulated genetic process for a 10-policy portfolio

Classical deterministic techniques like hill-climbing or steepest ascent can converge to a local optimum and become unable to escape unless an appropriate initial solution is chosen, as shown conceptually in the left panel in Figure 5. Evolutionary searches, on the other hand, due to their ability to "jump" through valleys in the objective function, are typically able to escape local maxima and continue the search of the optimization space, as shown in the right panel. Also, because genetic algorithms test the performance of different combinations of genes through many iterations, the space is searched efficiently. Combinations that result in poor performance are quickly pruned from the solution pool, and there is thus no need to scan all possible solutions. Consequently, a genetic algorithm can solve the 100-policy portfolio optimization problem in a matter of minutes with a standard desktop computer.

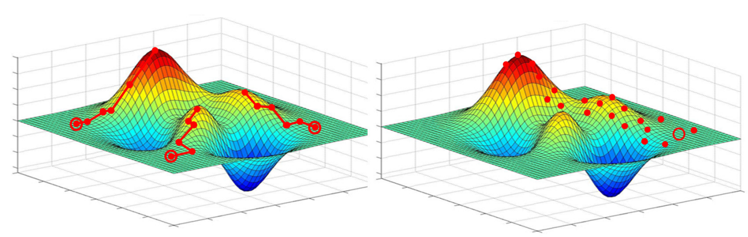

Figure 5. Conceptual comparison of deterministic hill-climbing techniques (left) vs. evolutionary searches (right)

CASE STUDY

A primary insurance company in Florida with a portfolio of properties distributed throughout the state wants to explore underwriting strategies that can help them attain several objectives, discussed below. Their total portfolio consists of more than 50,000 policies—1,000 of which are up for renewal next month (the remaining 49,000 policies are considered their static portfolio). They perform the following portfolio optimization exercises

Maximize Risk-Based Return (RBR)

In this first exercise, the company wishes to maximize their overall risk-based return (RBR) by selecting a subset of policies from the 1,000 that are up for renewal. They define RBR as total premium divided by the expected annual loss with an exceedance probability of 1%. The company uses three methodologies: policy ranking, steepest ascent, and AIR's evolutionary search algorithms. All three reach a similar solution, namely that the company should renew a subset of about 700 polices, while dropping the remaining ones that were responsible for decreasing the RBR into a state wind pool.

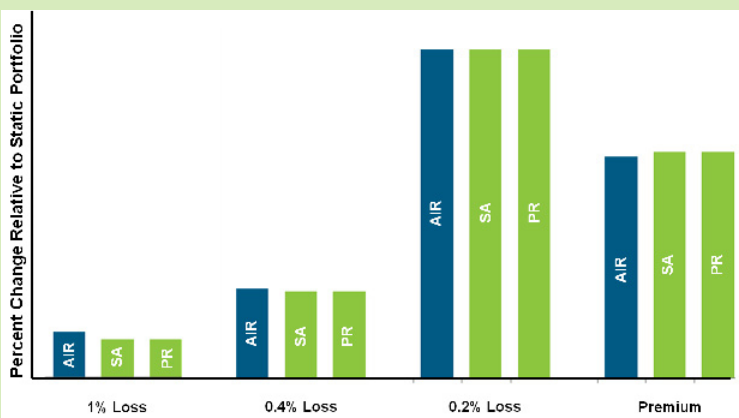

Figure 6. Maximizing risk-based return using three algorithms—AIR, steepest ascent (SA), and policy ranking (PR).

Maximize RBR While Limiting Expected Losses

In the second exercise, they again wish to maximize RBR, but while also constraining expected losses at the 1%, 0.4%, and 0.2% exceedance probability levels. This problem is more complex because the three constraints severely limit the potential combinations of policies, and the consideration of the loss correlation between the policies within a portfolio becomes a critical issue. The company applies the same three methodologies as before. Policy ranking and steepest ascent yield similar solutions, but AIR's evolutionary search algorithms deliver a solution that outperforms the others by yielding a 13% premium increase.

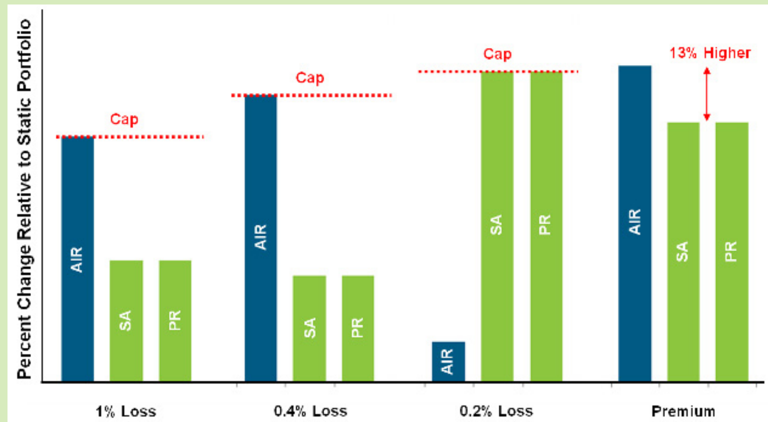

Figure 7. Maximizing risk-based return while limiting expected losses—AIR, steepest ascent (SA), and policy ranking (PR)

The underwriting team plots the locations of the policies and finds that the solutions from deterministic algorithms involve selecting policies that are geographically different than those from AIR's search. For instance, the SA and PR solutions would lead to dropping policies around the Tampa, Central Florida, and Northern Florida areas. In contrast, the solution obtained using AIR's algorithms suggests that keeping most of the policies in these areas while dropping some around the Miami-Dade area would yield a higher premium and keep the expected losses below the desired level

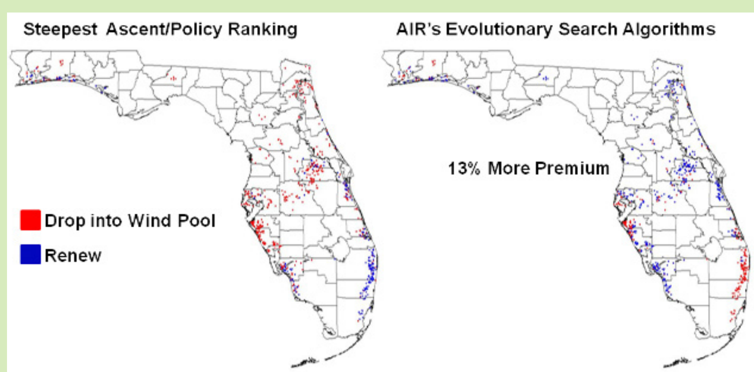

Figure 8. Comparison of policy locations based on the different algorithms

Maximize Premium and Limit Reinsurance Costs

In the third exercise, the company wishes to keep their reinsurance costs—which they calculate based on the tail value at risk (TVaR) of the entire portfolio—under a certain threshold. They optimize the portfolio using both the steepest ascent methodology and AIR's algorithms. They discover that AIR's algorithms provide a solution that yields 56% more premium than the deterministic algorithm solution, while keeping reinsurance costs at approximately the same level.

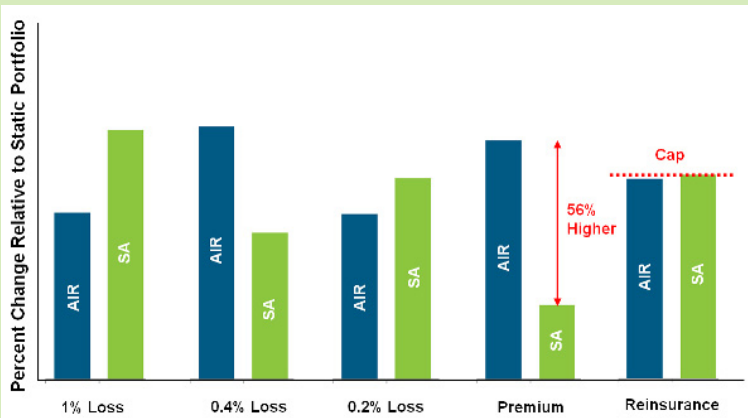

Figure 9. Maximizing premiums while limiting reinsurance costs—AIR and steepest ascent (SA)

Conclusion

The company in Florida, based on these numerical tests, concludes that while deterministic techniques may be suitable for handling simple problems with few or no constraints, the need to account for policy loss correlation in complex constrained scenarios requires advanced analytical custom solutions. They also discover that performance in underwriting and risk management can be enhanced with the help of computerized decision analytics.

AIR'S PORTFOLIO OPTIMIZATION SOLUTIONS

AIR has developed several techniques to approach optimization problems that were previously deemed too complex to be handled in a systematic, automated fashion. This allows portfolio managers to consider multiple business objectives when determining renewal strategies that accommodate their underwriting guidelines and their overall approach to enterprise risk management.

AIR currently offers portfolio optimization solutions on a consulting basis. Please contact us if you would like more information.

FURTHER READING

For more information on the application of evolutionary concepts to computation, refer to "Adaptation in Natural and Artificial Systems" by John Holland. For more details about genetic algorithms, refer to "Genetic Algorithms in Search, Optimization, and Machine Learning" by David E. Goldberg.

REFERENCES

[1]IT GOES WITHOUT SAYING THAT THE BIOLOGICAL PROCESSES THAT OCCUR IN NATURE ARE MUCH MORE COMPLEX THAN WHAT IS REPRESENTED IN THESE THREE STEPS. HOWEVER, PROGRAMMATICALLY, THESE MECHANISMS CAPTURE THE PROCESS SUFFICIENTLY WELL FOR OUR PURPOSES.

ABOUT AIR WORLDWIDE

AIR Worldwide (AIR) is the scientific leader and most respected provider of risk modeling software and consulting services. AIR founded the catastrophe modeling industry in 1987 and today models the risk from natural catastrophes and terrorism in more than 50 countries. More than 400 insurance, reinsurance, financial, corporate, and government clients rely on AIR software and services for catastrophe risk management, insurance-linked securities, detailed site-specific wind and seismic engineering analyses, agricultural risk management, and property replacement-cost valuation. AIR is a member of the Verisk Insurance Solutions group at Verisk Analytics and is headquartered in Boston with additional offices in North America, Europe, and Asia. For more information, please visit www. air-worldwide.com.

AIR